Damped Posterior Linearization Filter

Matti Raitoharju, Lennart Svensson, Ángel F. García-Fernández, and Robert Piché

arXiv:1704.01113v1 [math.OC] 4 Apr 2017

***Abstract*—The iterated posterior linearization filter (IPLF) is an algorithm for Bayesian state estimation that performs the measurement update using iterative statistical regression. The main result behind IPLF is that the posterior approximation is more accurate when the statistical regression of measurement function is done in the posterior instead of the prior as is done in non-iterative Kalman filter extensions. In IPLF, each iteration in principle gives a better posterior estimate to obtain a better statistical regression and more accurate posterior estimate in the next iteration. However, IPLF may diverge. IPLF's fixed-points are not described as solutions to an optimization problem, which makes it challenging to improve its convergence properties. In this letter, we introduce a double-loop version of IPLF, where the inner loop computes the posterior mean using an optimization algorithm. Simulation results are presented to show that the proposed algorithm has better convergence than IPLF and its accuracy is similar to or better than other state-of-the-art algorithms.**

***Index Terms*—Bayesian state estimation; nonlinear; estimation; Kalman filters**

I. Introduction

In Bayesian state estimation, a state that evolves stochastically in time is estimated from noisy measurements. In this letter, we concentrate on the measurement update stage of the Bayesian filter. In the measurement update, a prior (the dynamic model's state prediction) is updated using information from a measurement. If the measurement model is linear and Gaussian the posterior density can be computed analytically using the Kalman filter [1], but for a general measurement model, the computation of the posterior density is intractable. One approximate approach is the general Gaussian filter (GGF), which represents the joint state and measurement distribution by a multivariate Gaussian using moment matching, and then uses standard marginalization and conditioning formulas to compute the conditional state distribution given the measurement to obtain the posterior estimate [2]. The GGF moment matching is also intractable, but there exist several Kalman filter extensions (KFEs) that are approximations of the GGF. The standard way to apply GGF is to do statistical linear regression (SLR) in the prior, which can be interpreted as the optimal linearization with the given prior. The GGF approach does not work well for severe nonlinearities or low measurement noise [3] so there is interest in developing alternatives.

Iterative algorithms generate approximate solutions based on previous solutions so that the approximations are supposed to improve. The iterated extended Kalman filter (IEKF) [4, pp. 349-351] produces a sequence of mean estimates by making first-order Taylor approximations of the measurement function. In [5], it was shown that the IEKF measurement update is equivalent to the Gauss-Newton (GN) algorithm for computing the maximum a posteriori (MAP) estimate. IEKF performs well in situations where the true posterior is close to a Gaussian. Convergence of the GN algorithm is not guaranteed. Better convergence is offered by damped (or descending) GN algorithms [6], which ensure that the cost function is nonincreasing at each iteration. An alternative to damped GN is the Levenberg-Marquadt algorithm [7].

The iterated posterior linearization filter (IPLF) [8] uses SLR instead of the first order Taylor linearization. SLR's linearization is based on a larger area instead of a point as in the Taylor linearization. This improves the accuracy of the filter. In [8], it was shown that the desirable linearization would use the posterior instead of the prior. The basic idea in the IPLF is to compute a posterior estimate using the prior and then use this estimate to compute a better linearization and posterior estimate and so on. In this letter, we observe that the IPLF does not always converge and propose an adjusted version of the algorithm with improved convergence properties.

We show in simulations how the proposed algorithm is less prone to filter divergence than the original IPLF. Furthermore, we compare the posterior estimate accuracy with other iterative Kalman filter methods and show that the proposed algorithm is more accurate in the simulations.

II. Background

We consider measurements of the form

$$y = h(x) + \varepsilon, \tag{1}$$

where y is n_y dimensional real measurement value, $h(\cdot)$ is the measurement function, x is the n_x dimensional real random state vector and ε is zero mean Gaussian noise with covariance R.

A. GGF

The GGF [2] for measurement model (1) uses moments for predicted measurement $\hat{y}$, cross covariance of measurement and state $P_{h(x)x}$, and measurement covariance $P_{h(x)h(x)}$. These moments are

$$\hat{y} = \int h(x) p_{\mathrm{N}}(x; \mu^-, P^-) \mathrm{d}x \tag{2}$$

This article is based upon work from COST Action TU1302 (SAPPART), supported by COST (European Cooperation in Science and Technology).

Financial support by the Academy of Finland under grant no. #287792 (OpenKin) is hereby gratefully acknowledged.

M. Raitoharju and R. Piché are with the laboratory of Automation and Hydraulic Engineering, Tampere University of Technology, Finland. E-mail: first-name.last-name@tut.fi. L. Svensson is with the department of Signals and Systems, Chalmers University of Technology, Sweden. E-mail: lennart.svensson@chalmers.se. Á. F. García-Fernández is with the Department of Electrical Engineering and Automation, Aalto University, Finland. E-mail: angel.garciafernandez@aalto.fi

$$P_{h(x)x} = \int (h(x) - \hat{y}) \left(x - \mu^-\right)^T p_{\mathrm{N}}(x; \mu^-, P^-) \mathrm{d}x \quad (3)$$

$$P_{h(x)h(x)} = \int (h(x) - \hat{y}) (h(x) - \hat{y})^T p_{\mathrm{N}}(x; \mu^-, P^-) \mathrm{d}x, \quad (4)$$

where $p_{\mathrm{N}}(x|\mu^-, P^-)$ is the multivariate normal probability density function (PDF) of the prior that has mean μ^- and covariance P^-. The update equations for the mean and covariance of the posterior are then

$$\mu^+ = \mu^- + K(y - \hat{y}) \quad (5)$$

$$P^+ = P^- - KSK^T, \quad (6)$$

where

$$S = P_{h(x)h(x)} + R \quad (7)$$

$$K = P_{xh(x)} S^{-1}. \quad (8)$$

The integrals (36)-(4) do not have in general closed form solutions. Many KFEs, such as unscented Kalman filter (UKF) [9] and cubature Kalman filter (CKF) [10], can be interpreted as numerical approximations of GGF.

B. IPLF

In [8], it is shown that a lower bound of the Kullback-Leibler divergence (KLD) of a GGF with respect to true posterior is minimized if the moments (36)-(4) are computed using the posterior instead of the prior. This update cannot be computed directly with the GGF equations (36)-(8), but instead a linear model that has correct moments is generated and is then applied to the prior.

The IPLF algorithm iteratively approximates the posterior using SLR, creating a sequence of updated means μ_i and covariance matrices P_i as follows. First, set $\mu_1 = \mu^-$ and $P_1 = P^-$. Then, at the ith iteration, compute the SLR of $h(\cdot)$ with respect to μ_i, P_i by first computing moments (36)-(4) and then defining the linear measurement that corresponds to those moments [8]:

$$\hat{h}(x) = J_i x + b_i + \varepsilon_{\Omega i} + \varepsilon, \quad (9)$$

where

$$J_i = {P_{xh(x)}}_i^T P_i^{-1} \quad (10)$$

$$b_i = \hat{y}_i - J_i \mu_{i-1} \quad (11)$$

$$\Omega_i = {P_{h(x)h(x)}}_i - J_i P_i J_i^T \quad (12)$$

$$\varepsilon_{\Omega i} \sim \mathrm{N}(0, \Omega_i), \quad (13)$$

$\varepsilon_{\Omega i}$ is an independent noise whose covariance Ω_i is the covariance of the linearization error. For linear systems Ω_i is 0.

Finally, compute the posterior mean and covariance that correspond to the linearized measurement function (9):

$$S = J_i P^- J_i^T + R + \Omega_i \quad (14)$$

$$K_i = P^- S^{-1} \quad (15)$$

$$\mu_{i+1} = \mu^- + K_i(y - J_i \mu^- - b_i) \quad (16)$$

$$P_{i+1} = P^- - K_i S_i K_i^T. \quad (17)$$

TABLE I
MEAN ESTIMATES OF FIRST 6 ITERATIONS OF IPLF AND IEKF

Iteration	1	2	3	4	5	6
IPLF	-2.51	5.28	-31.75	17.52	-40.56	11.96
IEKF	-7.64	58.28	-1.77	2.60	-6.66	48.47

The obtained posterior is used for the next linearization. The process is repeated for a predetermined number of steps or until a convergence criterion is met, e.g. until the KLD of two consecutive estimates is below a threshold.

Next we show an example where IPLF diverges:

Example 1: The measurement model is

$$y = \arctan(x) + \varepsilon. \quad (18)$$

State x has prior mean 2.75 and variance 1, measurement value is $y = 0$ and measurement noise variance is $R = 10^{-4}$. The integrals (36)-(4) are computed using Monte Carlo (MC) integration with 10^5 samples and also with IEKF.

Table I gives the means in the first 6 iterations. Neither of the algorithms converge to the true posterior, which is close to 0, even after 50 iterations.

C. IEKF as a GN algorithm

IEKF is similar to IPLF, but it uses first order Taylor approximations of the measurement function at a point rather than statistical linear regression. The IEKF iteration formulas are as in Section II-B, but with (39)-(12) replaced by

$$J_i = \left.\frac{\mathrm{d}h(x)}{\mathrm{d}x}\right|_{\mu_{i-1}} \quad (19)$$

$$b_i = h(\mu_{i-1}) - J_i \mu_i \quad (20)$$

$$\Omega_i = 0. \quad (21)$$

In [5], it is shown that the IEKF for measurements of form (1) is equal to the GN algorithm [6] that minimizes the cost function

$$\begin{aligned} q(\mu) =& \frac{1}{2}(h(\mu) - y)^T R^{-1} (h(\mu) - y) \\ &+ \frac{1}{2}(\mu - \mu^-)^T P^{-1} (\mu - \mu^-). \end{aligned} \quad (22)$$

To decrease the possibility of the divergence of GN, the damped (or descending) version of the GN algorithm can be used [6]. A damped IEKF that uses line search [6] in minimization of (22) is given in [11]. The mean update is

$$\mu_{i+1} = (1 - \alpha)\mu_i + \alpha(\mu_0 + K_i(y - h(\mu_i))) \quad (23)$$

with step length parameter α selected so that $0 \leq \alpha \leq 1$ and $q(\mu_{i+1}) \leq q(\mu_i)$. Finally, in the IEKF, the posterior covariance matrix is computed based on the linearization at the MAP point after the last iteration [5].

III. DAMPED ITERATED POSTERIOR LINEARIZATION FILTER (DIPLF)

The main weakness with the IPLF is that it sometimes diverges. In this section, we present a damped version of the IPLF, which has substantially better convergence properties.

The key insight is that the update of the mean can be seen as an optimization problem and solved with a damped GN algorithm that decreases the cost function in every iteration. A detailed motivation of the algorithm is presented in Section III-A, followed by a description of a set of techniques that are needed to implement it in Section III-B.

A. Derivation of DIPLF

Ideally, we want the posterior mean and covariance to be used in the SLR to linearize the measurement function [8]. Our approach to iteratively finding the mean and covariance for the linearization is to use a separate loop for the optimization of the mean. Thus, the proposed filter computes the estimate using two nested loops. The inner loop uses a damped GN algorithm to optimize the mean μ_i while keeping P_j and Ω_j fixed. The outer loop updates the covariances P_j and Ω_j. Index i increases in every iteration of the inner loop and j in the outer loop.

1) Inner loop: The purpose of the inner loop is to find optimal μ_i when P_j and Ω_j are fixed. The posterior mean μ that coincides with the mean used in the SLR is given by

$$\mu_i = \mu^- + K(\mu_i)(y - J(\mu_i)\mu^- - b(\mu_i)), \tag{24}$$

where we used (μ_i) to identify vectors and matrices that depend on the mean. In the Appendix, we show that the above equation is fulfilled when function

$$\begin{aligned} q(\mu_i) =& \frac{1}{2}(\hat{y}(\mu_i) - y)^T (R + \Omega_j)^{-1} (\hat{y}(\mu_i) - y) \\ &+ \frac{1}{2}(\mu_i - \mu^-)^T P^{-^-1}(\mu_i - \mu^-)) \end{aligned} \tag{25}$$

achieves its minimum. Therefore, we use it as the cost function for the inner loop. The optimization problem is analogous to the optimization of the cost function of the IEKF (22), but using $\hat{y}$ instead of $h(\mu)$. To ensure that $q(\cdot)$ does not increase, we use a damping factor α ($0 < \alpha \leq 1$) in the update as in (23)

$$\mu_{i+1} = (1 - \alpha)\mu_i + \alpha\left(\mu_0 + K_i(y - J_i\mu_0 - b_i)\right) \tag{26}$$

and select an α such that the value of (25) decreases. This makes the algorithm locally convergent on almost all nonlinear least squares problems, provided that the line search is carried out appropriately. In fact, it is usually globally convergent [6, Chapter 9.2.1]. We propose that the algorithm should repeat the inner loop as long as (25) reduces significantly at each iteration i.e. while

$$q(\mu_i, P_j) < \beta q(\mu_i, P_j) \tag{27}$$

with, e.g., $\beta = 0.9$.

2) Outer loop: The outer loop updates the covariances P_j and Ω_j for the next iteration of the inner loop. The outer loop requires a stopping condition. To define this stopping condition, we first exponentiate function (25) and normalize to get a product of normal distributions:

$$\begin{aligned} &\frac{1}{\sqrt{(2\pi)^{n_y}|R + \Omega_j|(2\pi)^{n_x}|P^-|}} e^{-q(x)} \\ =& p_{\mathrm{N}}(\hat{y}_i | y, R + \Omega_j) p_{\mathrm{N}}(\mu_i | \mu^-, P^-). \end{aligned} \tag{28}$$

Algorithm 1: DIPLF

```
i ← 0, j ← 0, μ_0 ← μ^-, P_0 ← P^-, Ω_0 ← 0
Compute J_i and ŷ_i at prior using (36), (3), and (39)
while p_N(ŷ_i|y, R + Ω_j) p_N(μ_i|μ^-, P^-) increases do
    j ← j + 1
    Compute P_j using (17)
    Compute Ω_j using (12)
    while q(μ_i, P_j) < β q(μ_{i-1}, P_j) do
        i ← i + 1
        Find 0 < α_i ≤ 1 for (26) so that (25) becomes
          smaller using line search (Section III-B3)
        Compute μ_i using (26)
        Compute J_i using (3) and (39)
        Compute ŷ_i using (36)
    end
end
Use the mean and covariance from the second last round
  of the outer loop, i.e. those that had highest
  p_N(ŷ_i|y, R + Ω_j) p_N(μ_i|μ^-, P^-), as the posterior
  estimate
```

If we treat P_i and Ω_i as constants, (25) achieves its minimum when (28) achieves its maximum. We thus monitor that (28) is increasing as the stopping condition of the outer loop. Compared to using (25), the factor containing Ω_i in (28) causes the stopping condition (28) to favor estimates with smaller nonlinearity.

3) Full algorithm: The algorithm is summarized in Algorithm 1. The inner loop does the optimization of the target function (25) while keeping the state covariance and nonlinearity covariance constant. The direction of change in a GN algorithm is a descent direction of the cost function [6, Chapter 9.2.1]. Thus, with small enough α the cost function value should decrease, unless the estimate is in a local extremum or a saddle point. By using $\alpha = 1$ and exiting the inner loop after each iteration, Algorithm 1 becomes equivalent to an IPLF algorithm.

The outer loop changes the optimization problem of the inner loop as Ω changes and the covariance, for which the expected measurement value $\hat{y}$ is computed, changes.

B. Implementation considerations

1) Implementation with sigma-point filters: Integrals (36)-(4) do not have closed form solutions in general. Sigma-point filtering is a common framework for approximating these integrals [2]. The formulas are

$$\hat{y} \approx \sum_{k=1}^{m} w_k h(\mu + \Delta_k) \tag{29}$$

$$P_{h(x)x} \approx \sum_{k=1}^{m} w_k \left(h(\mu + \Delta_k) - \hat{y}\right) \Delta_k^T \tag{30}$$

$$P_{h(x)h(x)} \approx \sum_{k=1}^{m} w_k \left(h(\mu + \Delta_k) - \hat{y}\right) \left(h(\mu + \Delta_k) - \hat{y}\right)^T, \tag{31}$$

TABLE II
KLDs OF ESTIMATES PRODUCED BY DIFFERENT FILTERS IN arctan TEST

	GGF	RUF	IPLF	DIPLF
MC	15.9	0.04	88.69	$3 \cdot 10^{-6}$
EKF	4009.10	0.01	65.12	10^{-6}
CKF	3370.78	0.01	64.39	10^{-6}
UKF	92.55	0.01	10^{-6}	10^{-6}

TABLE III
MEAN KLDs OF ESTIMATES PRODUCED BY DIFFERENT FILTERS IN RANGE TEST

	GGF	KLPUKF	RUF	IPLF	DIPLF
MC	0.25	0.17	0.21	0.26	0.17
EKF	0.48	0.60	0.48	0.55	0.55
CKF	0.28	0.22	0.29	0.38	0.23
UKF	0.35	0.34	0.33	0.37	0.26

where w_k is the weight of kth sigma-point, m is the number of sigma points, and Δ_k is the translation of kth sigma-point from the state mean. The translation vectors Δ_k depend on the state dimension and the state covariance matrix P and their computation involves taking a matrix square root (e.g. Cholesky decomposition) of P. In Algorithm 1 the state covariance matrix is constant in the inner loop and so the Δ_k vectors need to be recomputed only in the outer loop.

2) Efficient computation of the Jacobian: The updated state covariance (17) can be written in information form using the matrix inversion lemma:

$$P_{j+1} = \left(J_i^T (R + \Omega_j)^{-1} J_i + P^{--1} \right)^{-1}. \tag{32}$$

Using this in (39) we get

$$J_{i+1} = P_{h(x)x} \left(J_i^T (R + \Omega_j)^{-1} J_i + P^{--1} \right). \tag{33}$$

Since P^- does not change between iterations, P^{--1} needs to be computed only once for the inner loop and matrices that have to be inverted at every iteration of the inner loop have dimension $n_y \times n_y$. Thus, when $n_y \ll n_x$ (33) is faster than (39).

3) Line search: The value of α can be computed with different line search techniques [6, Chapter 9.2.1]. Backtracking line search, also known as the Armijo-Goldstein step length principle, is one of the simplest and most widely used options, and is done as follows.

First, α is set to 1, and a value of μ is found using (23). If this value does not decrease (25), w.r.t. the previous value of μ, then α is multiplied by a constant factor $0 < \tau < 1$ and the procedure is repeated.

IV. SIMULATIONS

In our first test, we evaluate the situation in Example 1. We compare posterior estimates obtained with estimation algorithms that are based on approximating integrals (36)-(4). We use MC integration with 10^5 samples, extended Kalman filter (EKF), UKF with parameters from [9], and third-degree CKF [10] to compute the moment approximations. The algorithms used are GGF, recursive update filter (RUF) [12], IPLF [8], and DIPLF. The GGF with different moment computation algorithms is equivalent to the algorithms from which the moment computation was taken. RUF applies the measurement using down-weighted Kalman gains and in our tests uses 10 iterations. When using EKF to compute the moments GGF is EKF, RUF is equivalent to the algorithm in [13], DIPLF is IEKF, and DIPLF becomes damped IEKF. We use $\tau = 0.5$ for reducing the step length. For the threshold of inner loop we use $\beta = 0.9$.

Results are given in Table II. DIPLF gives good posterior estimates with all moment computation methods. IPLF converges within 50 iterations only when using UKF for moment computation. We also found that whether IPLF converges is sensitive to used prior and moment computation algorithm. So changing the prior mean a little may make it to converge with different moment computation algorithms. RUF has slightly worse estimates than DIPLF and GGF does not provide good estimates with any of the tested moment computation methods.

In our second test, we computed estimates for two dimensional positioning using three range measurements. The prior was in origin with covariance I. Ranges were computed to sources at $\begin{bmatrix} -1 & 0 \end{bmatrix}^T$, $\begin{bmatrix} 0 & 1 \end{bmatrix}^T$ and $\begin{bmatrix} 1 & -2 \end{bmatrix}^T$. Measurement noise covariance was I and measurements were generated by sampling a true location from the prior and then generating corresponding measurements. In this test we used the same filters as in the previous test and Kullback-Leibler partitioned update Kalman filter (KLPUKF) [14]. KLPUKF uses a linear transformation to decorrelate measurement elements in such a way that the nonlinearity of a measurement element is minimized and applies measurement elements sequentially. For the test with a single measurement, it would have been identical to GGF.

The test is repeated 1000 times using different true location and measurement values. The mean KLDs are presented in Table III. Results show that the iterative algorithms that use EKF moment computation, except RUF, produce worse estimates than EKF itself. This is probably caused by that the posterior may be multimodal in this test. With other moment computation methods the IPLF is worse than GGF, but damped iterated posterior linearization filter (DIPLF) is better than either of those. In this test, KLPUKF and DIPLF were the most accurate algorithms, except when tested with EKF.

V. CONCLUSIONS AND FUTURE WORK

We have shown that by fixing the covariance matrices of the IPLF algorithm, the optimal posterior mean in the posterior linearization approach is the solution of an optimization problem. We solved this optimization problem using a damped GN algorithm. Then we updated the covariance matrices P_j and Ω_j and repeated the mean optimization. Compared to IPLF this algorithm is less prone to diverge. The proposed algorithm was more accurate than other tested algorithms in our simulations.

In future, one could look how the posterior linearization could be done with other optimization algorithms than GN, such as Levenberg-Marquadt. Another interesting research topic would be extending the algorithm to work with non-additive noise.

REFERENCES

[1] Y. Ho and R. Lee, "A Bayesian approach to problems in stochastic estimation and control," *IEEE Transactions on Automatic Control*, vol. 9, no. 4, pp. 333 – 339, October 1964, doi:10.1109/TAC.1964.1105763.

[2] S. Särkkä, *Bayesian filtering and smoothing*. Cambridge University Press, 2013, vol. 3. [Online]. Available: http://users.aalto.fi/~ssarkka/pub/cup_book_online_20131111.pdf

[3] M. Morelande and Á.F. García-Fernández, "Analysis of Kalman filter approximations for nonlinear measurements," *IEEE Transactions on Signal Processing*, vol. 61, no. 22, pp. 5477–5484, Nov 2013, doi:10.1109/TSP.2013.2279367.

[4] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*, ser. Mathematics in Science and Engineering. Academic Press, 1970, vol. 64.

[5] B. Bell and F. Cathey, "The iterated Kalman filter update as a Gauss-Newton method," *Automatic Control, IEEE Transactions on*, vol. 38, no. 2, pp. 294–297, 1993.

[6] Å. Björck, *Numerical Methods for Least Squares Problems*. Society for Industrial and Applied Mathematics, 1996, doi:10.1137/1.9781611971484.

[7] R. L. Bellaire, E. W. Kamen, and S. M. Zabin, "New nonlinear iterated filter with applications to target tracking," vol. 2561, 1995, pp. 240–251, doi:10.1117/12.217701.

[8] Á.F. García-Fernández, L. Svensson, M. Morelande, and S. Särkka, "Posterior linearization filter: Principles and implementation using sigma points," *IEEE Transactions on Signal Processing*, vol. 63, no. 20, pp. 5561–5573, Oct 2015, doi:10.1109/TSP.2015.2454485.

[9] E. Wan and R. Van der Merwe, "The unscented Kalman filter for nonlinear estimation," in *Proceedings of the Adaptive Systems for Signal Processing, Communications, and Control Symposium. AS-SPCC.*, 2000, pp. 153–158, doi:10.1109/ASSPCC.2000.882463.

[10] I. Arasaratnam and S. Haykin, "Cubature Kalman filters," *IEEE Transactions on Automatic Control*, vol. 54, no. 6, pp. 1254–1269, June 2009, doi:10.1109/TAC.2009.2019800.

[11] M. A. Skoglund, G. Hendeby, and D. Axehill, "Extended Kalman filter modifications based on an optimization view point," in *Information Fusion (Fusion), 2015 18th International Conference on*, July 2015, pp. 1856–1861.

[12] Y. Huang, Y. Zhang, N. Li, and L. Zhao, "Design of sigma-point Kalman filter with recursive updated measurement," *Circuits, Systems, and Signal Process.*, pp. 1–16, August 2015, doi:10.1007/s00034-015-0137-y.

[13] R. Zanetti, "Recursive update filtering for nonlinear estimation," *IEEE Transactions on Automatic Control*, vol. 57, no. 6, pp. 1481–1490, June 2012, doi:10.1109/TAC.2011.2178334.

[14] M. Raitoharju, Á. F. García-Fernández, and R. Piché, "Kullback–Leibler divergence approach to partitioned update Kalman filter," *Signal Processing*, vol. 130, pp. 289 – 298, 2017, doi:10.1016/j.sigpro.2016.07.007.

APPENDIX

Here we derive the cost function for updating the inner loop. First, we modify the Kalman update similarly as in [5] to get a GN update from IEKF update, but taking Ω_j term of IPLF into account:

$$\begin{aligned}
\mu_i =&\mu^- + K(\mu_i)(y - J(\mu_i)\mu^- - b(\mu_i)) \\
=&\mu^- + (J(\mu_i)^T(R+\Omega_j)^{-1}J(\mu_i) + P^{-^{-1}})^{-1}J(\mu_i)^T(R+\Omega_j)^{-1}(y - J(\mu_i)\mu^- - \hat{y}(\mu_i) + J(\mu_i)\mu_i) \\
=&(J(\mu_i)^T(R+\Omega_j)^{-1}J(\mu_i) + P^{-^{-1}})^{-1}(J(\mu_i)^T(R+\Omega_j)^{-1}(y - \hat{y}(\mu_i) + J(\mu_i)\mu_i) + P^{-^{-1}}\mu^-) \\
=&(J(\mu_i)^T(R+\Omega_j)^{-1}J(\mu_i) + P^{-^{-1}})^{-1}(J(\mu_i)^T(R+\Omega_j)^{-1}(y - \hat{y}(\mu_i) + J(\mu_i)\mu_i) + P^{-^{-1}}\mu_i - P^{-^{-1}}\mu_i + P^{-^{-1}}\mu^-) \\
=&\mu_i + (J(\mu_i)^T(R+\Omega_j)^{-1}J(\mu_i) + P^{-^{-1}})^{-1}(J(\mu_i)^T(R+\Omega_j)^{-1}(y - \hat{y}(\mu_i)) + P^{-^{-1}}(\mu^- - \mu_i)).
\end{aligned} \tag{34}$$

From the last row we can see that the IPLF update has a stationary point when

$$J(\mu_i)^T(R+\Omega_j)^{-1}(\hat{y}(\mu_i) - y) + P^{-^{-1}}(\mu_i - \mu^-) = 0. \tag{35}$$

Differentiating the expected value of a measurement

$$\hat{y}(\mu_i) = \int h(x) p_{\mathrm{N}}(x; \mu_i, P_j)\mathrm{d}x \tag{36}$$

with respect to the mean μ_i gives

$$\begin{aligned}
&\frac{\mathrm{d}\int h(x)p(x;\mu_i,P_j)\mathrm{d}x}{\mathrm{d}\mu_i} \\
=&\int h(x)\frac{\mathrm{d}p(x;\mu_i,P_j)}{\mathrm{d}\mu_i}\mathrm{d}x \\
=&\int h(x)(x-\mu_i)^T p(x;\mu_i,P_j)\mathrm{d}x P_j^{-1}.
\end{aligned} \tag{37}$$

Noting that $\int \hat{y}(\mu_i)(x-\mu_i)^T p(x;\mu_i,P_j)\mathrm{d}x = 0$ we can write

$$\int (h(x) - \hat{y}(\mu_i))(x-\mu_i)^T p(x;\mu_i,P_j)\mathrm{d}x P_j^{-1} = P_{h(x)x}P_j^{-1}. \tag{38}$$

From this we see that the Jacobian computed from SLR

$$J(\mu_i) = {P_{xh(x)}}_i^T P_j^{-1} \tag{39}$$

is the Jacobian of the expected value of the measurement i.e. $J(\mu_i) = \frac{\mathrm{d}\hat{y}(\mu_i)}{\mathrm{d}\mu_i}$. Using this we see that the left hand side of (35) is the derivative of the function:

$$q(\mu_i) = \frac{1}{2}(\hat{y}(\mu_i) - y)^T(R+\Omega_j)^{-1}(\hat{y}(\mu_i) - y) + \frac{1}{2}(\mu_i - \mu^-)^T P^{-^{-1}}(\mu_i - \mu^-), \tag{40}$$

which has a local extremum when (35) holds. Thus, (40) can be used as the cost function of the inner loop.